



Building Consistent and Highly Available Distributed Systems with Apache Ignite and GridGain



Valentin Kulichenko

GridGain Lead Architect & Apache Ignite PMC

Agenda


- What is Ignite/GridGain?
- Ignite/GridGain and CAP Theorem
- High Availability in Ignite/GridGain
 - Cache Backups
 - Persistence
 - Data Center Replication
 - Data Snapshots (+Demo)
- Q&A



the **in-memory** computing platform
that is **durable**, **strongly consistent** and **highly available**
with powerful **SQL**, **key-value** and **processing** APIs



GridGain In-Memory Computing Platform



Financial
Services




Telco




Travel &
Logistics



E-Commerce



Pharma &
Healthcare



IoT

SQL

Key/Value

Transactions

Compute

Services

Streaming

ML

Memory-Centric Storage

Ignite Native Persistence
(Flash, SSD, Intel 3D XPoint)

Third-Party Persistence
(RDBMS, HDFS, NoSQL)



Data Snapshots & Recovery



Monitoring & Management

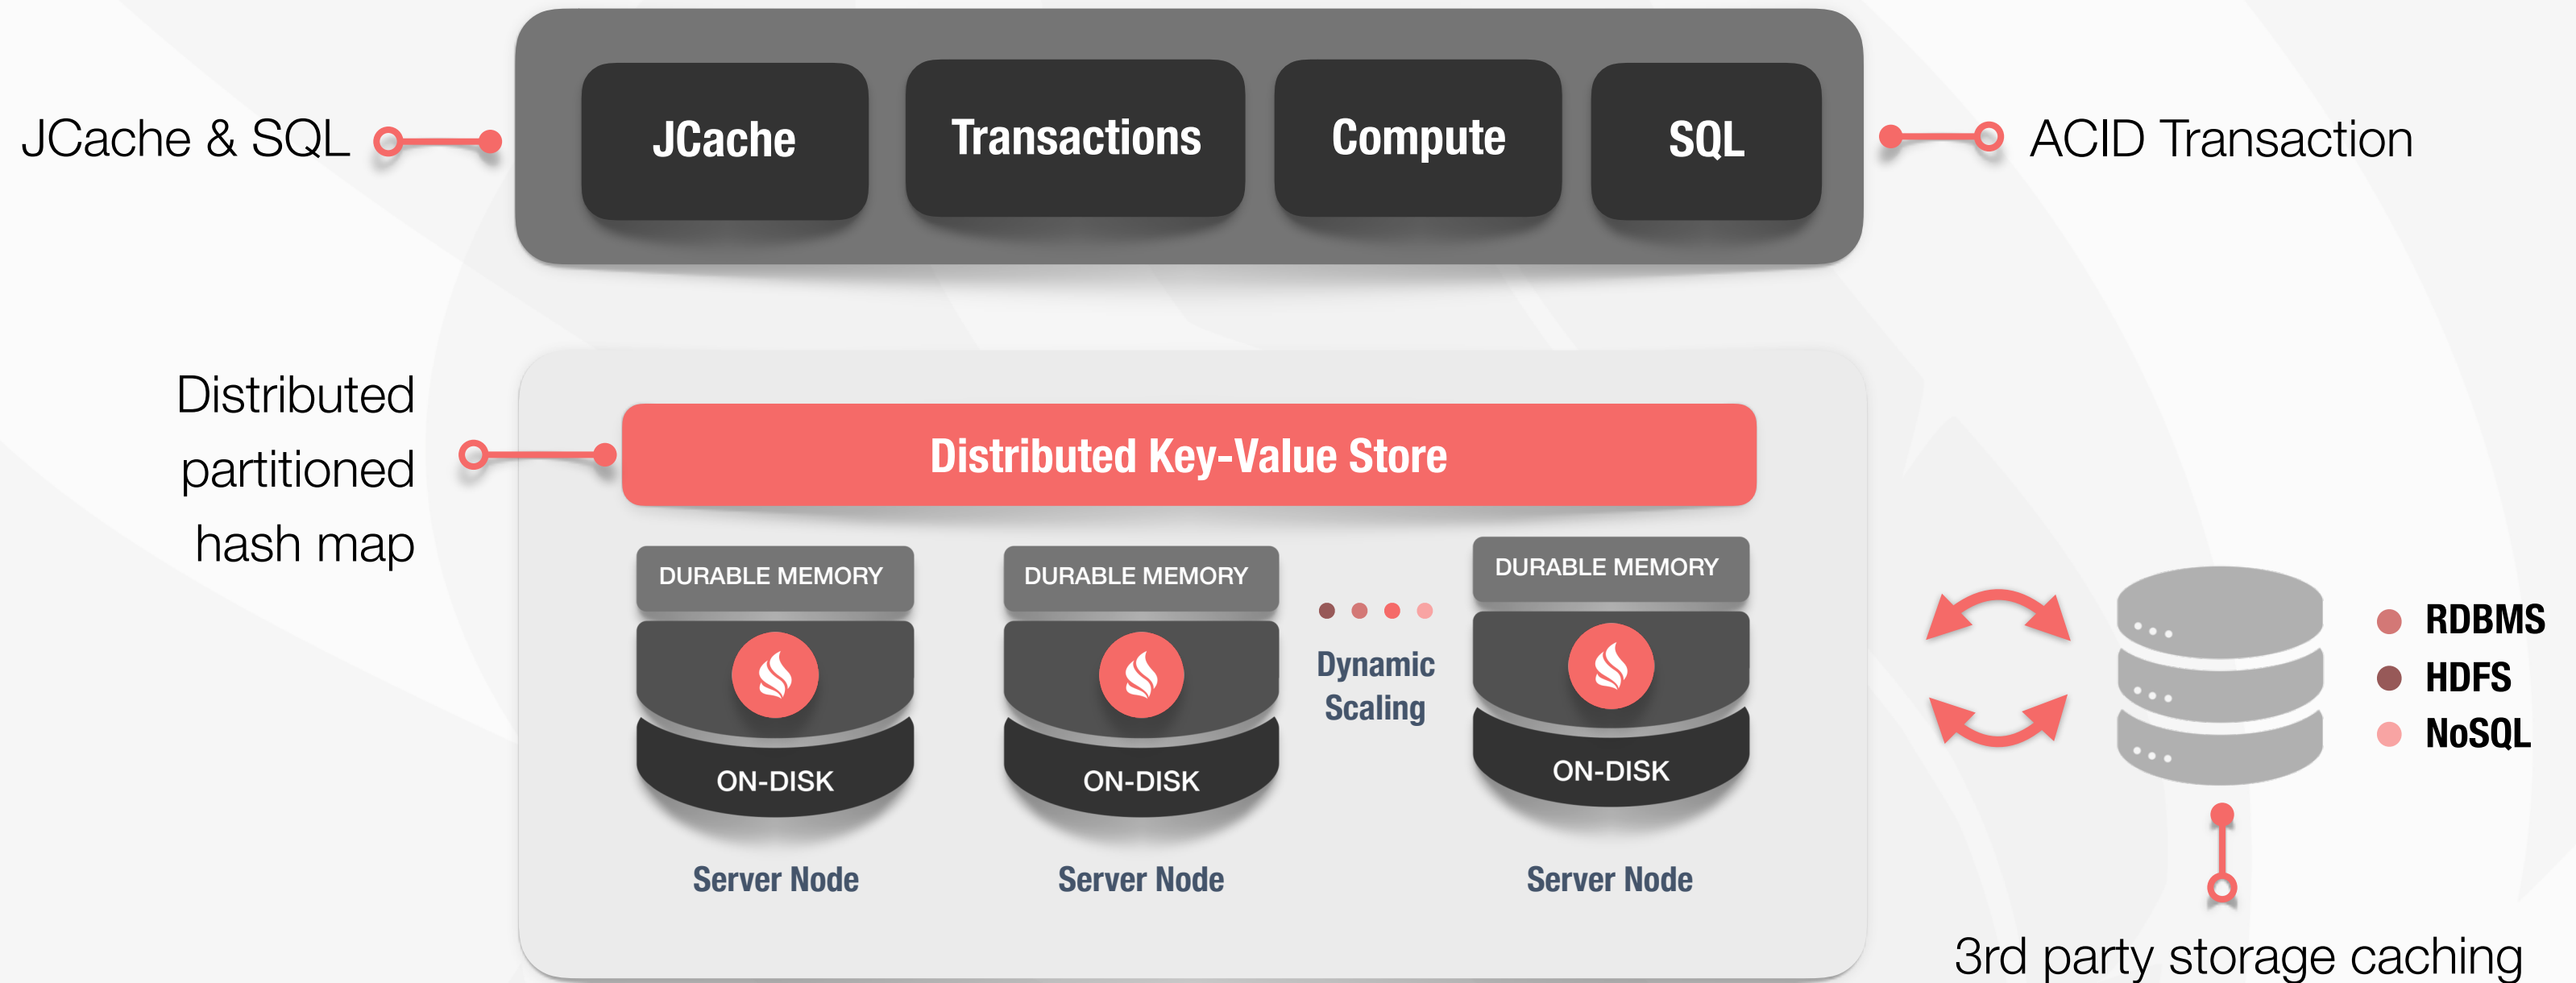


Security & Auditing

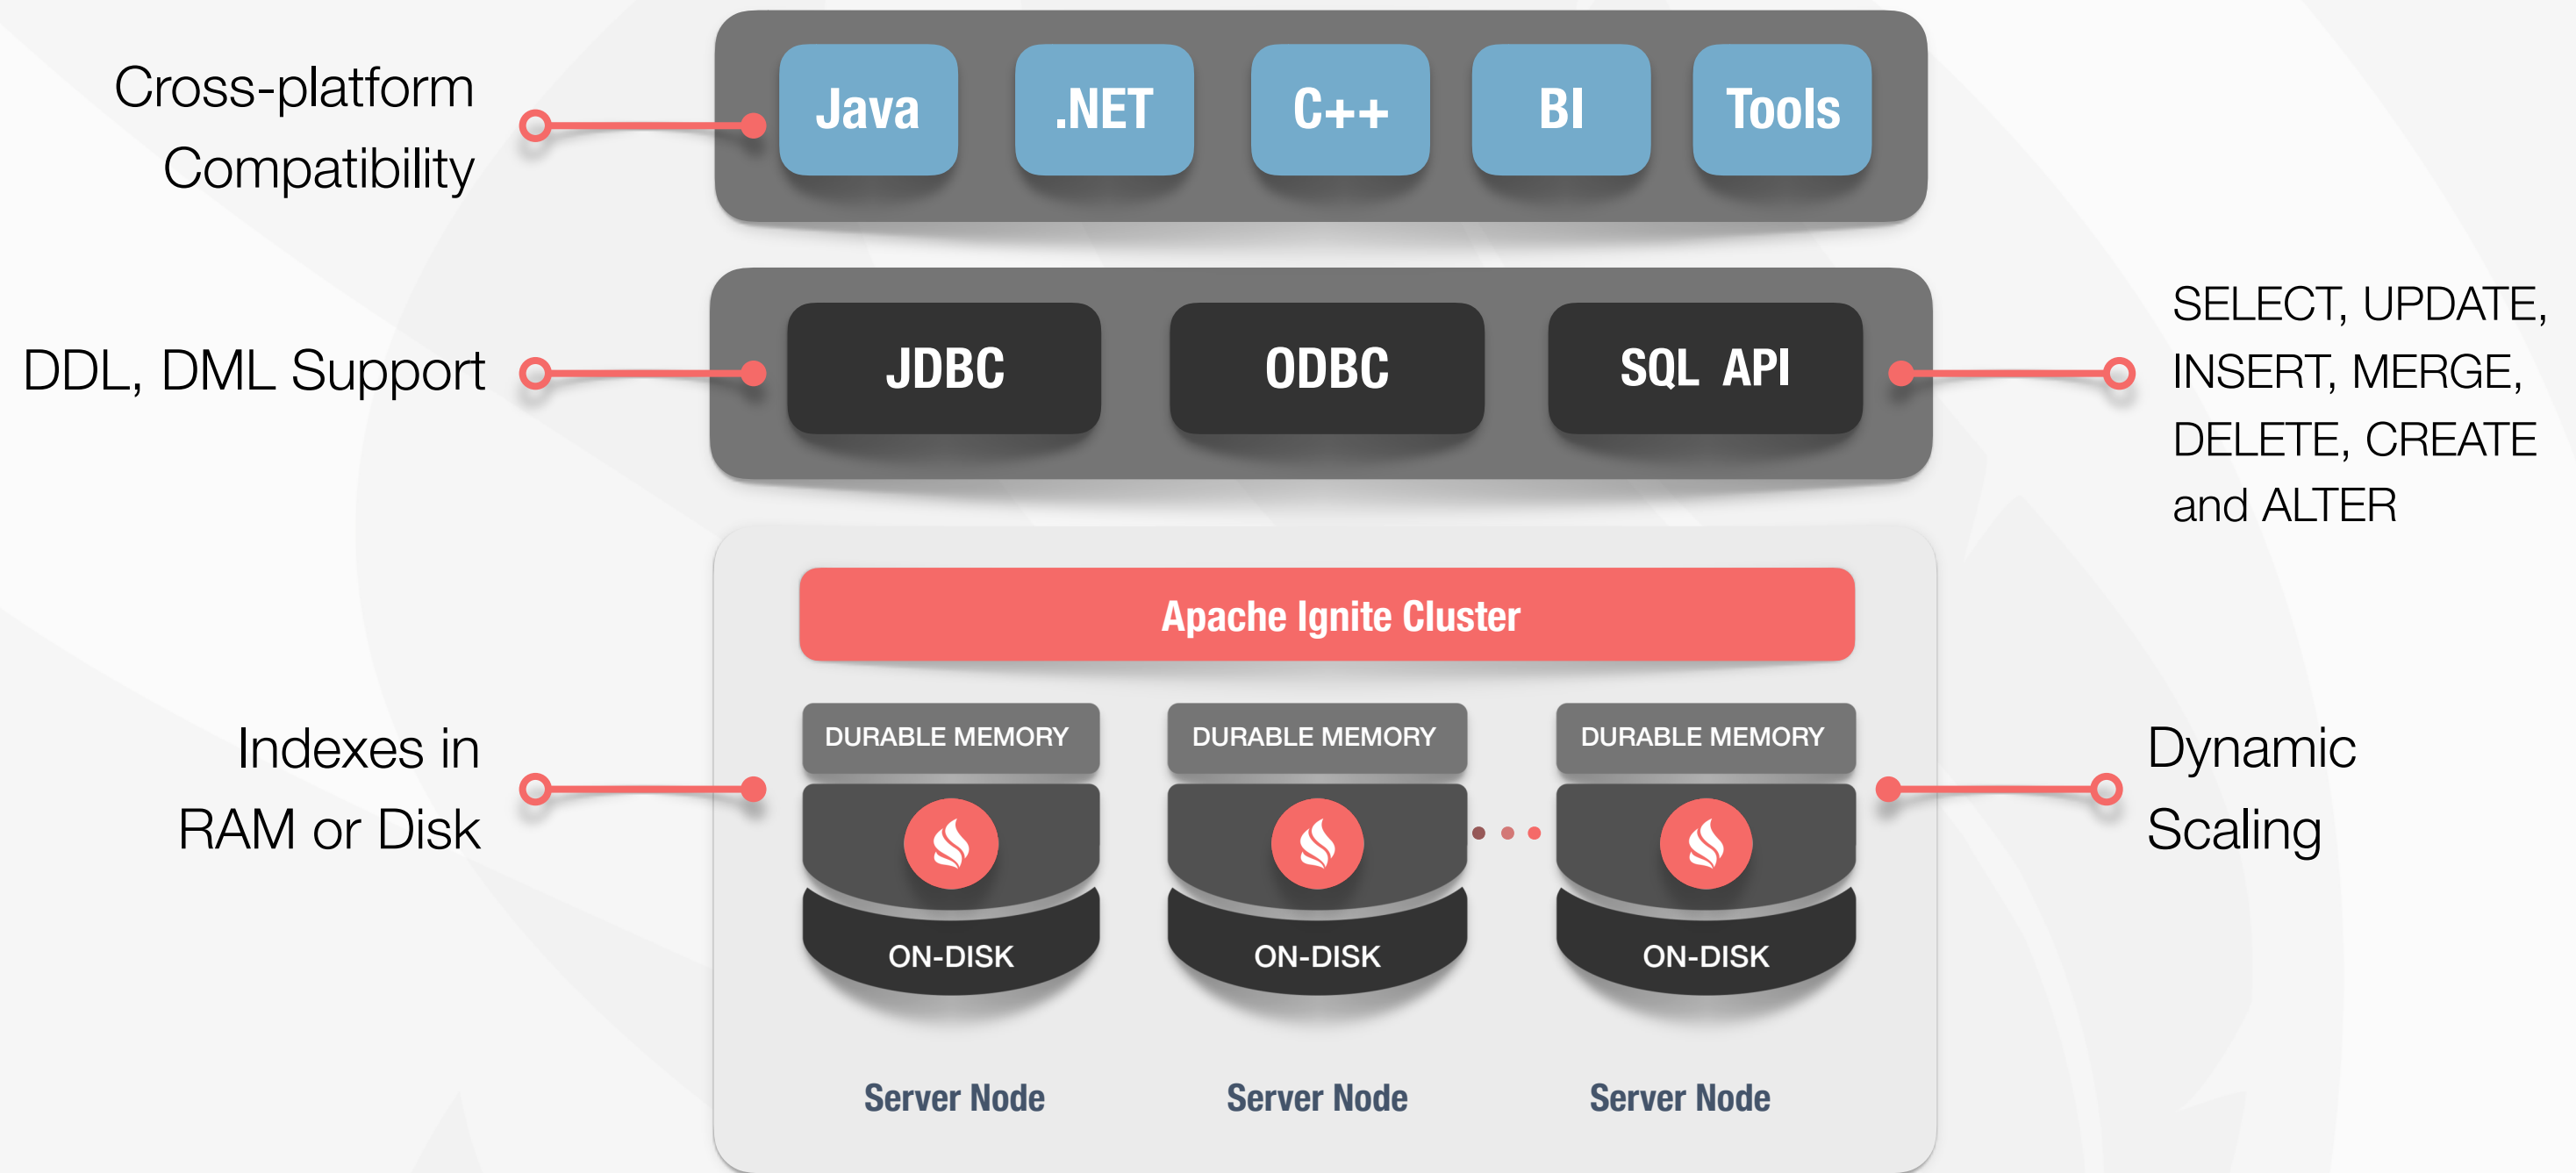


Data Center Replication

Data Grid



Distributed SQL

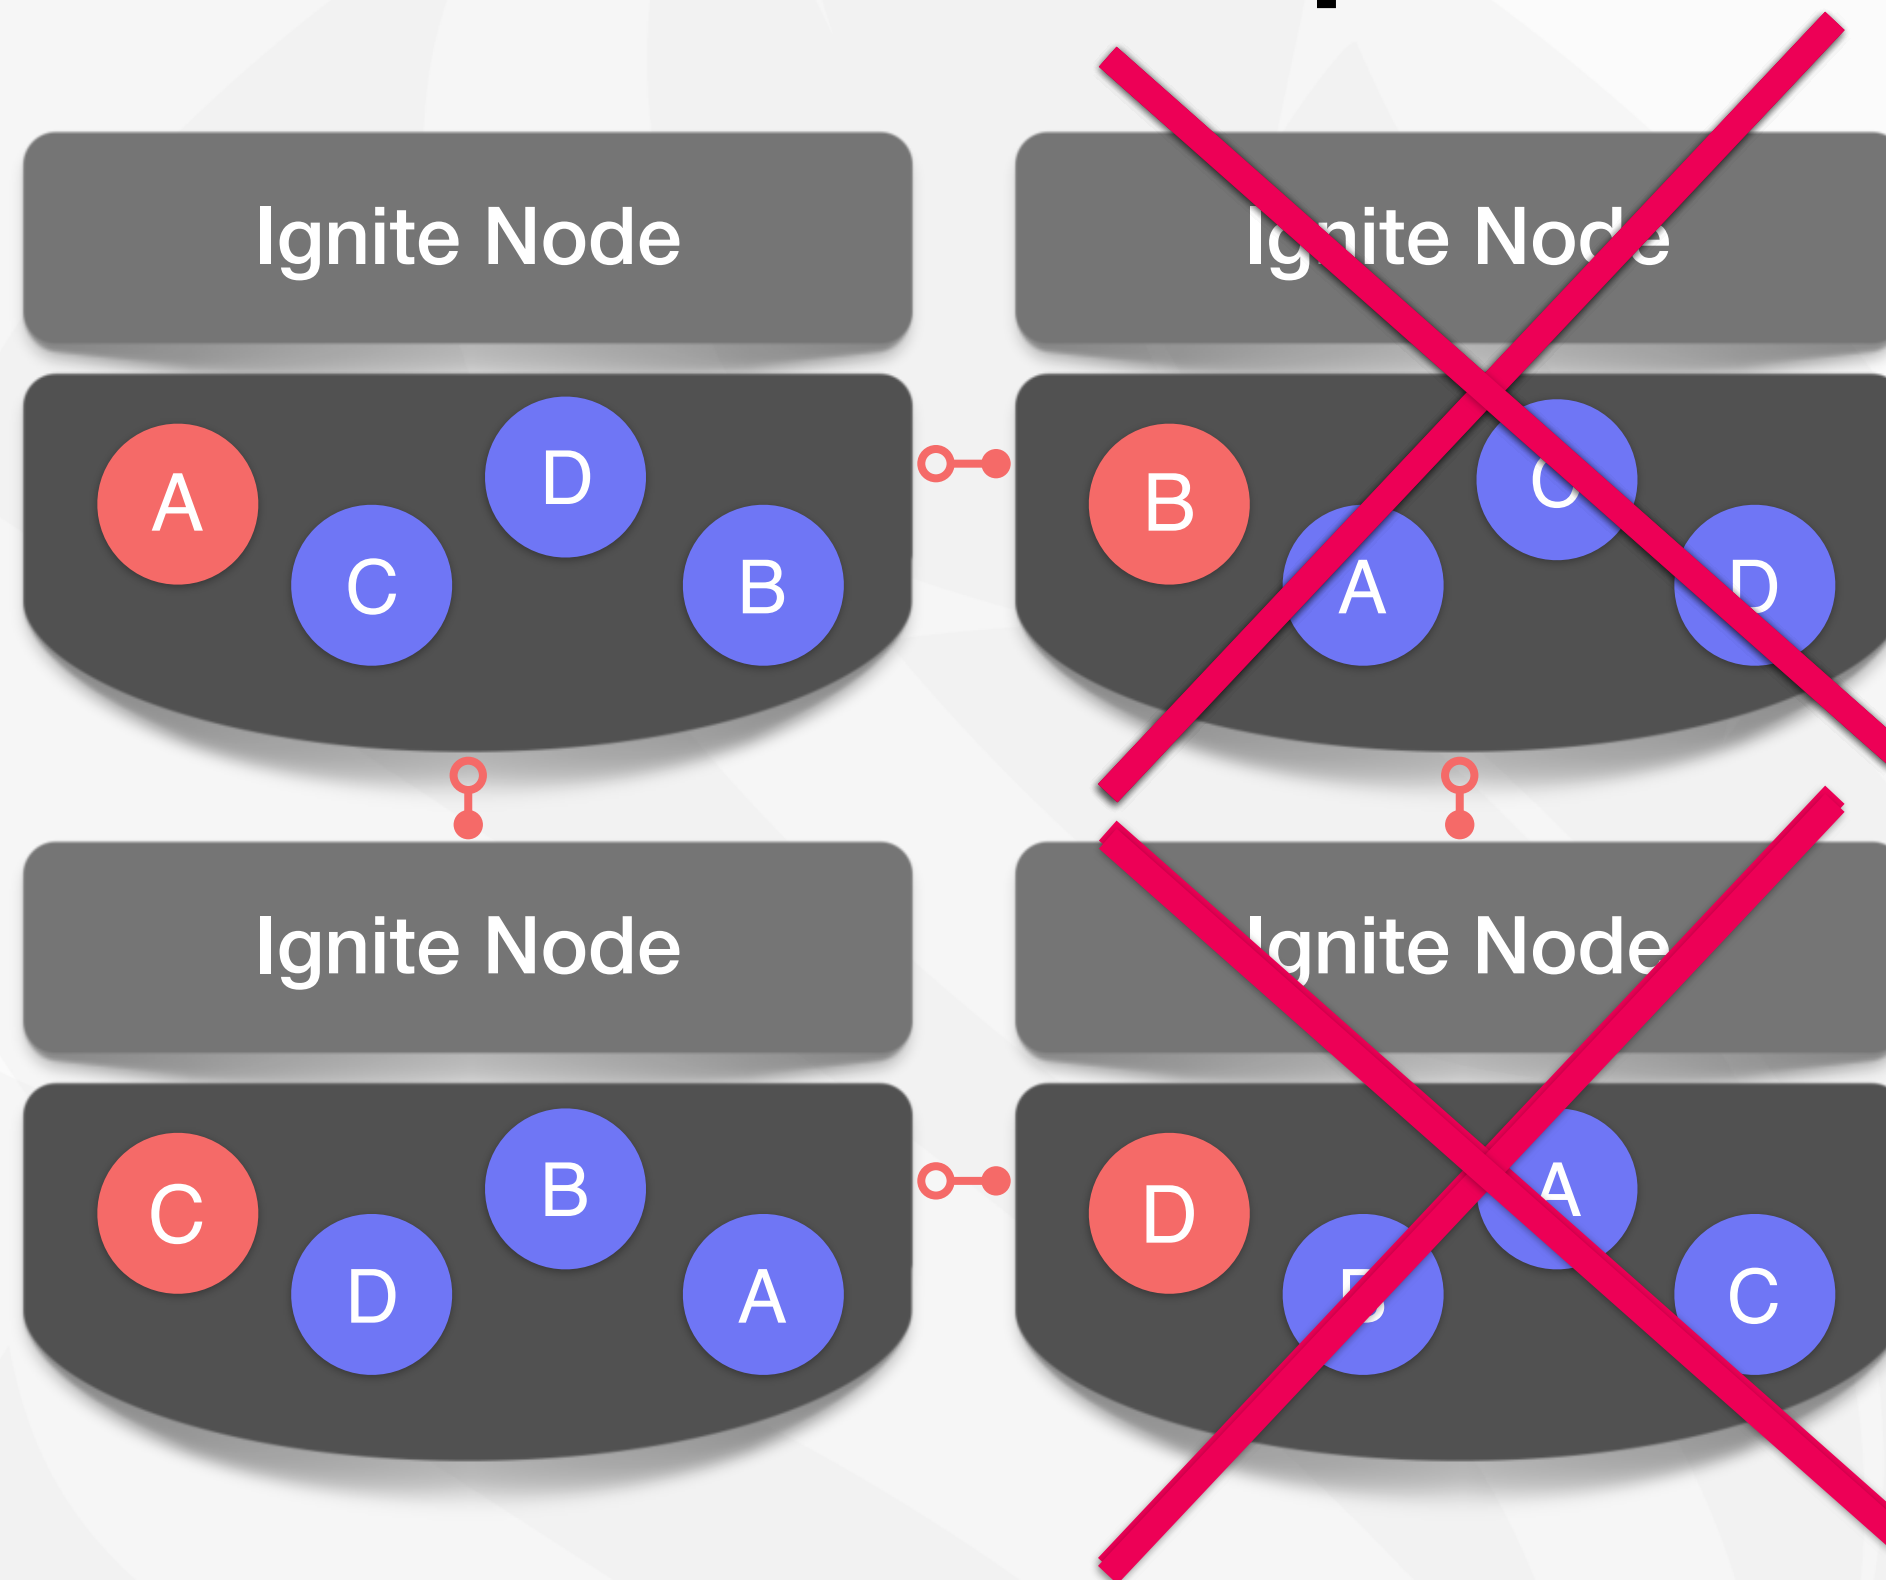


GridGain and CAP Theorem

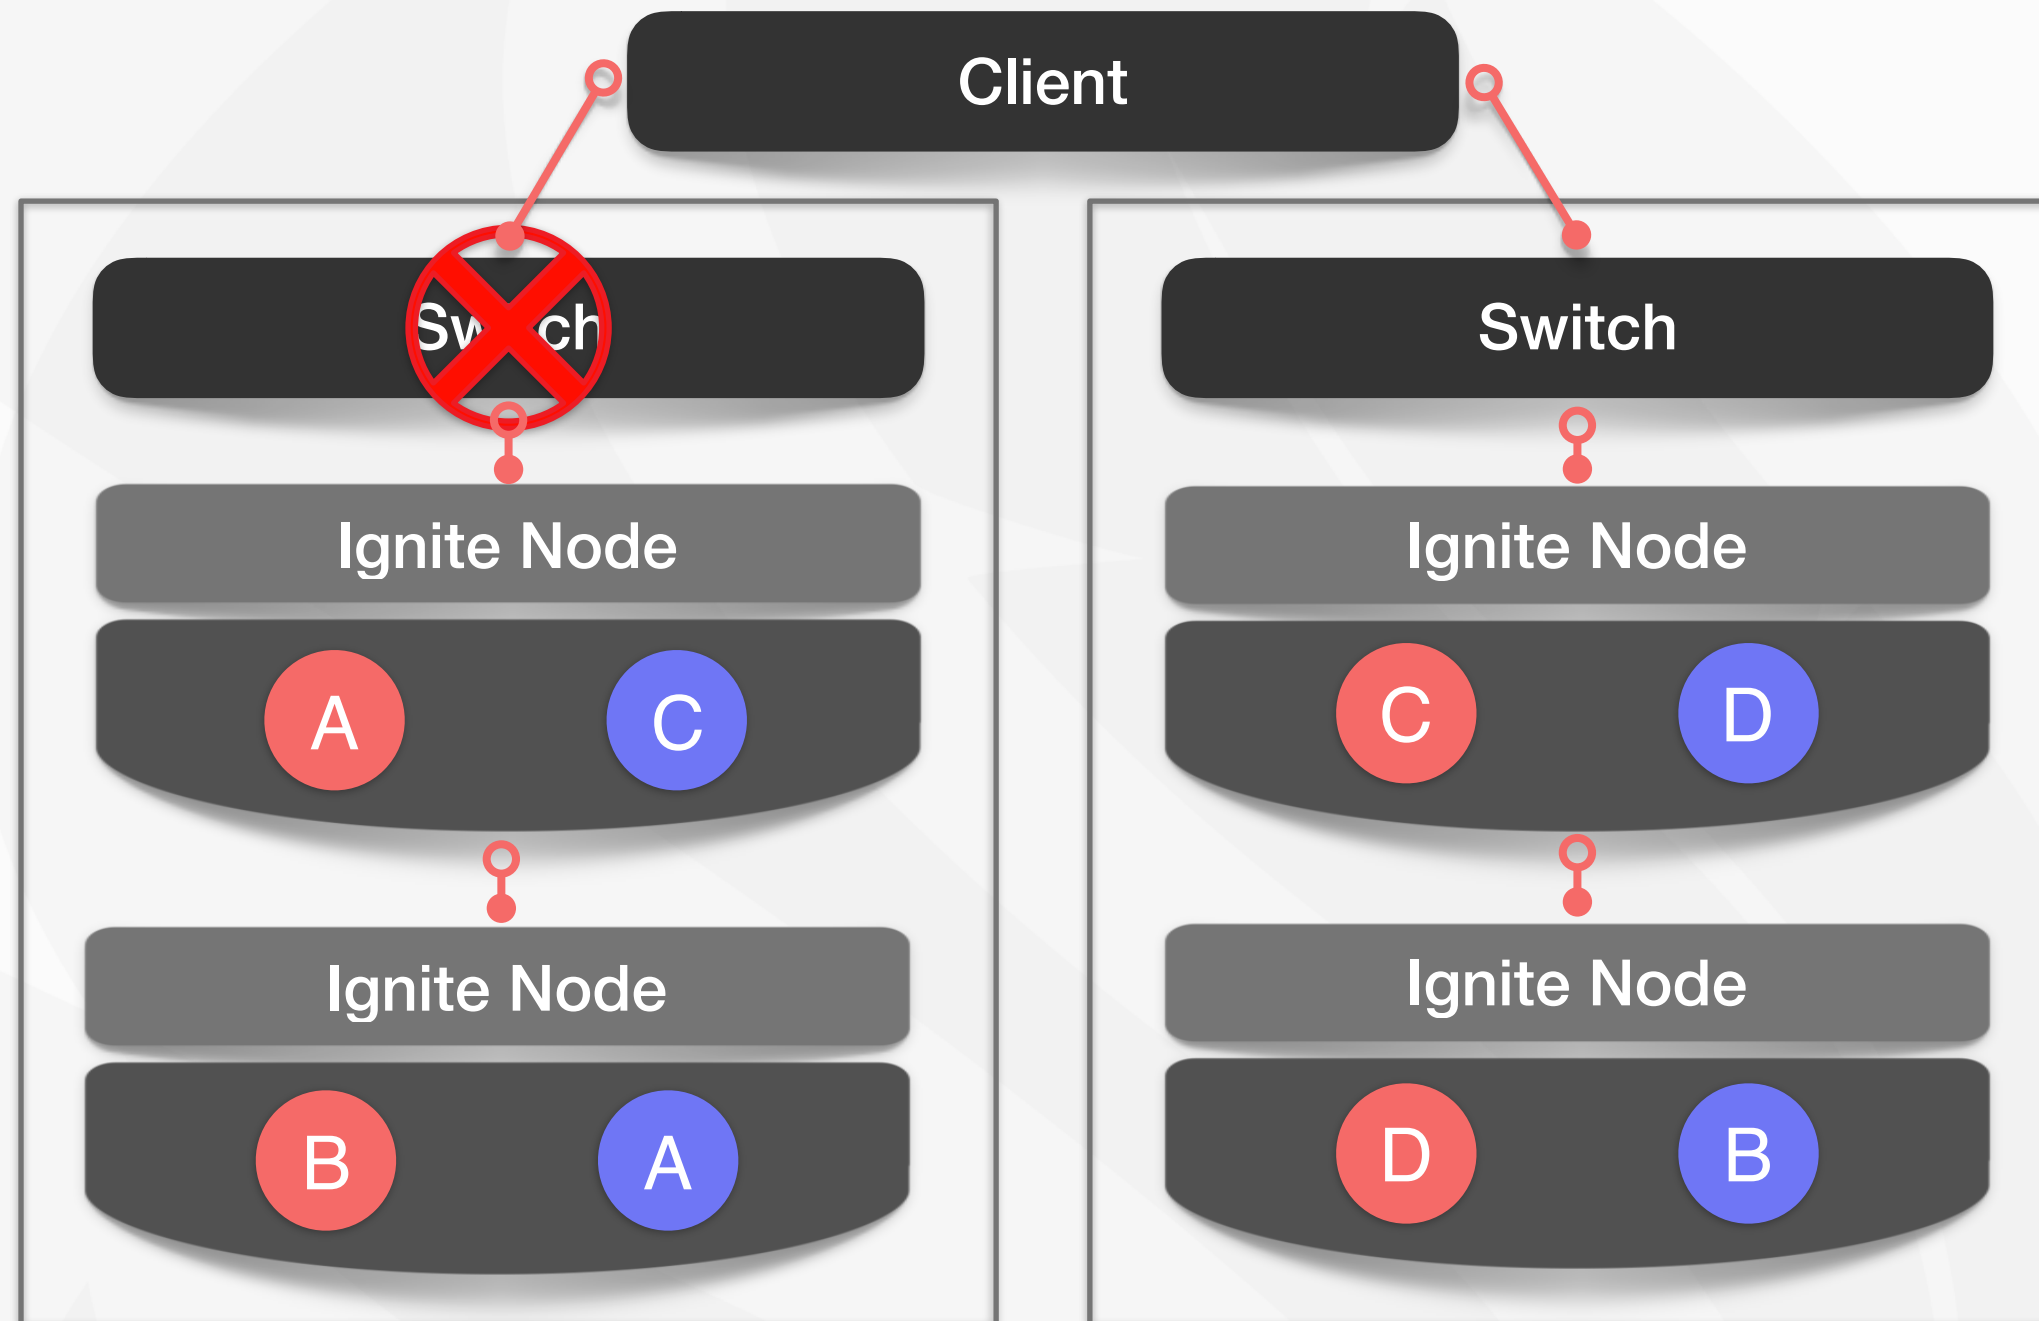
- CAP
 - Consistency (C)
 - Availability (A)
 - Partition Tolerance (P)
 - Impossible: CAP
 - Most of distributed systems are CP or AP
- GridGain
 - Strongly CP
 - But still highly available?

C
A?
P

Cache Backups



Rack Safety: Problem



Partition A is lost!

Rack Safety: Solution

- Assign attribute to each node:

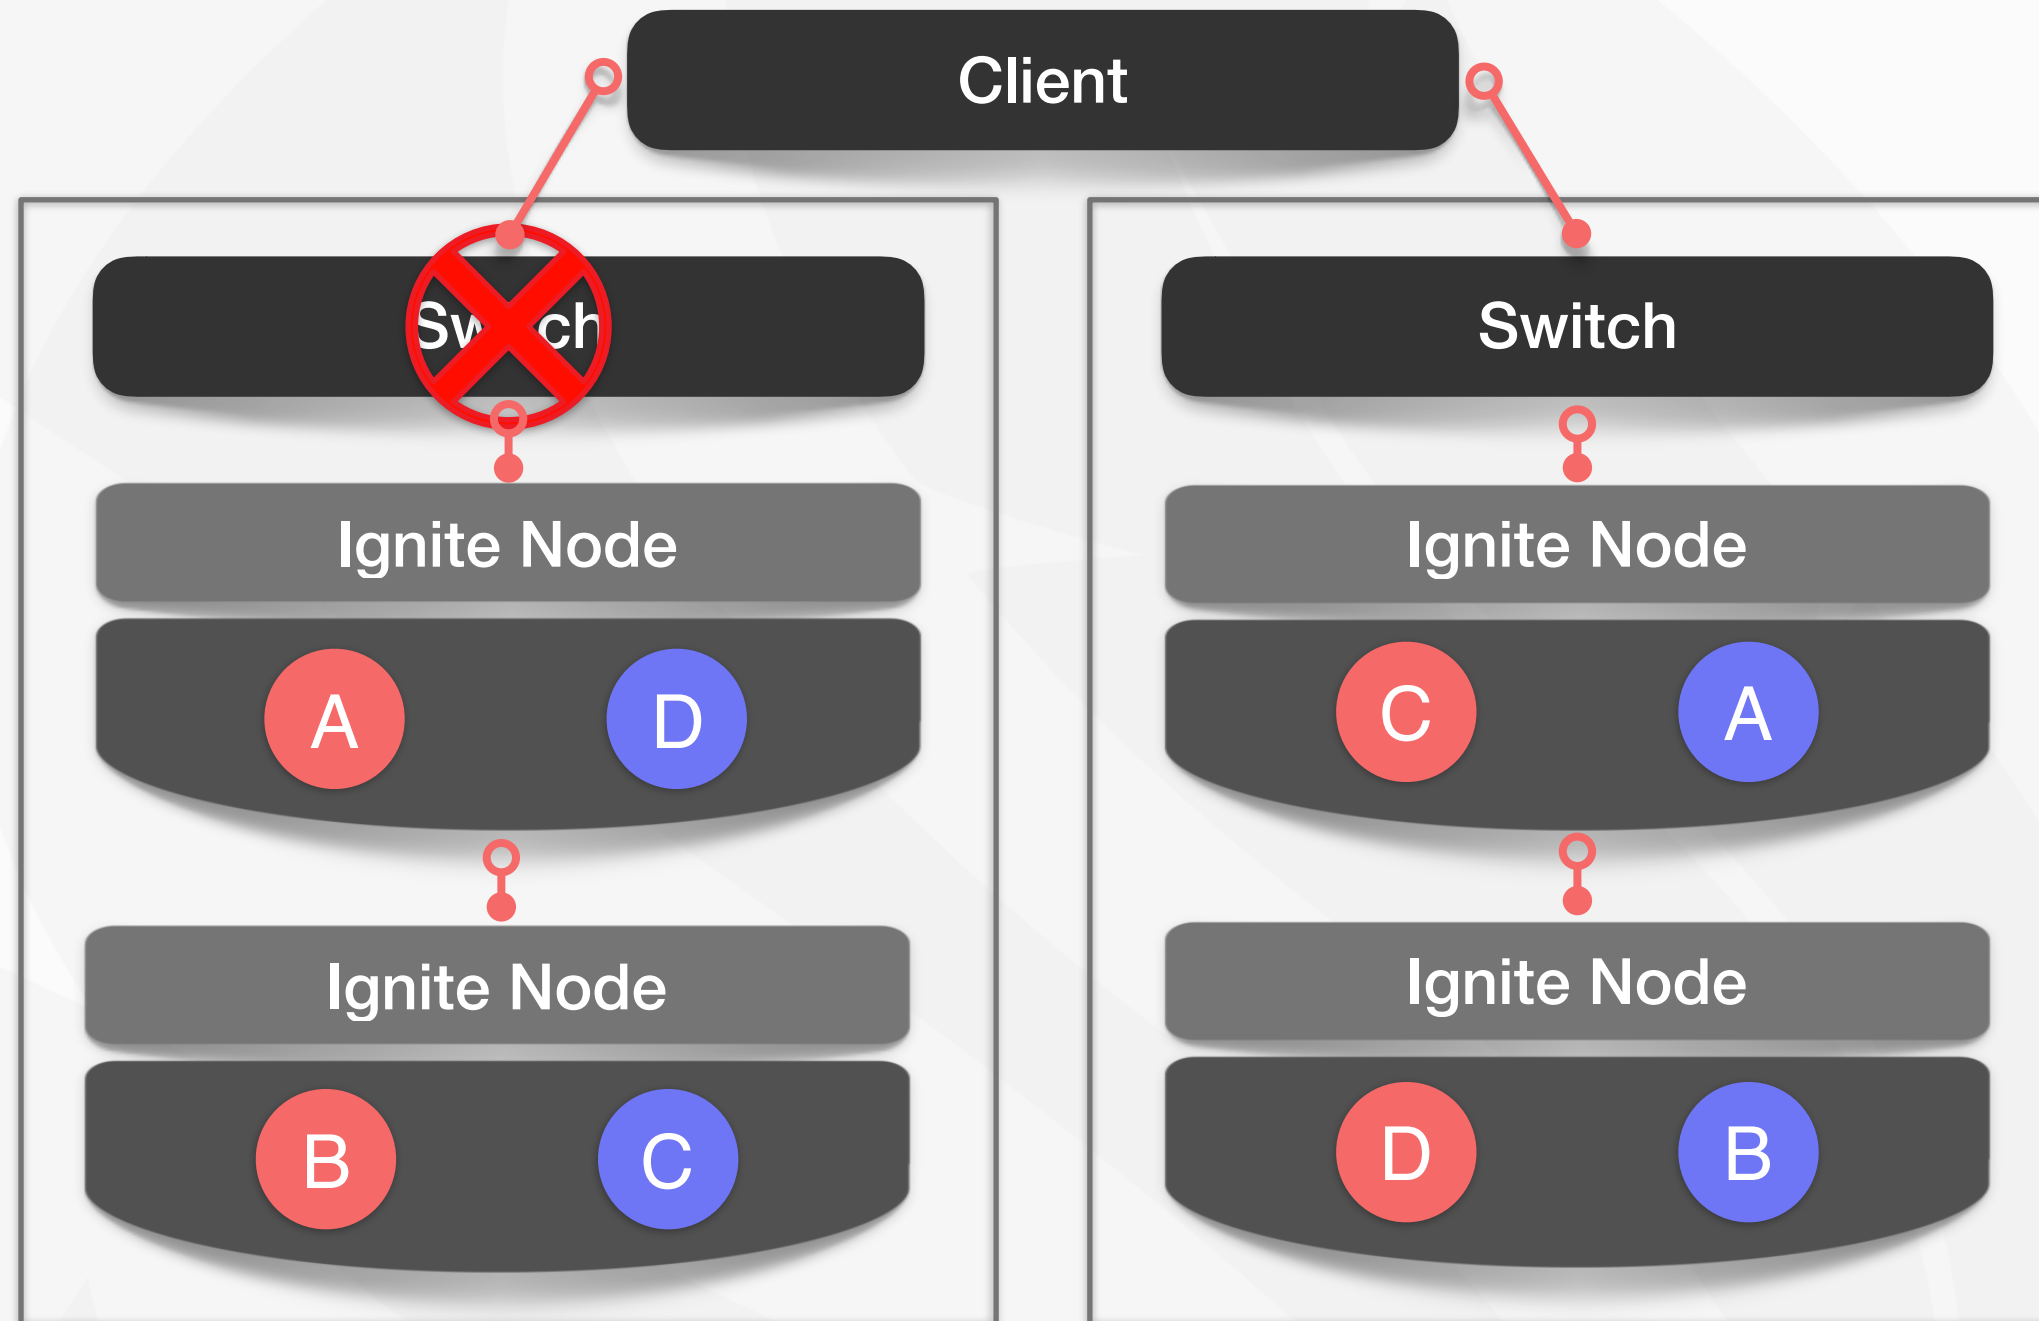
```
<property name="userAttributes">
  <map>
    <entry key="rack" value="rack-1"/>
  </map>
</property>
```

- Implement affinity backup filter:

```
public class RackFilter implements IgniteBiPredicate<ClusterNode, List<ClusterNode>> {
    @Override public boolean apply(ClusterNode candidate, List<ClusterNode> assigned) {
        String candidateRack = candidate.attribute("rack");
        String primaryRack = assigned.get(0).attribute("rack");

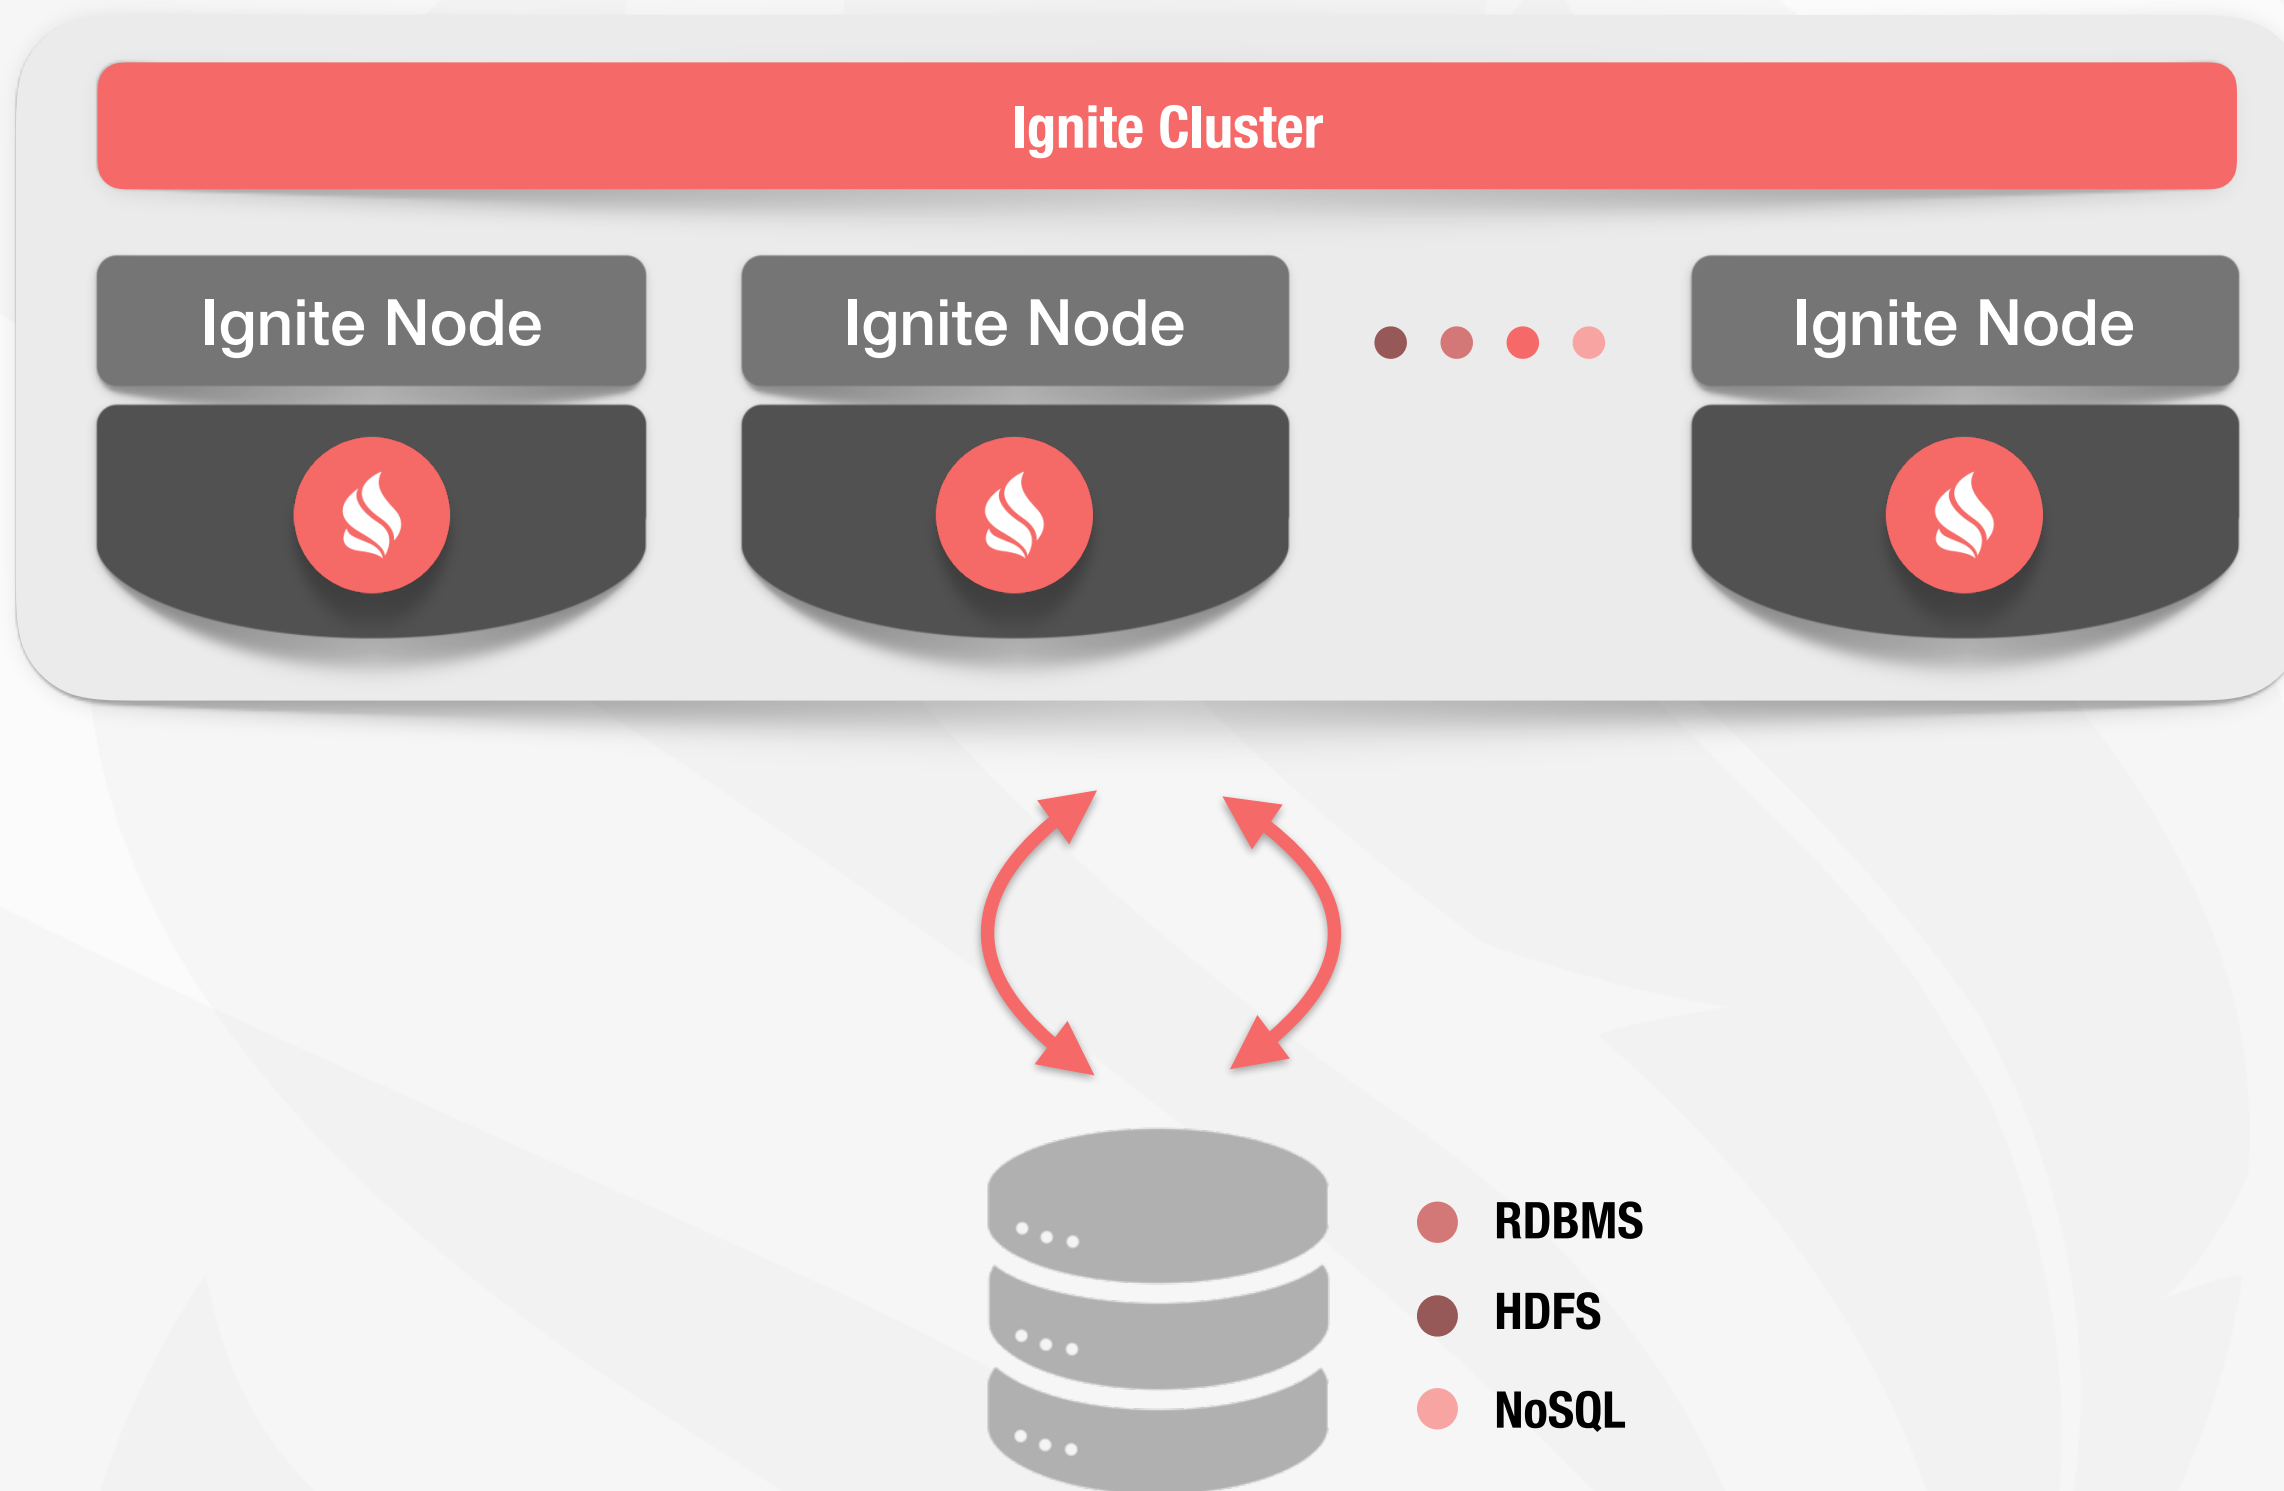
        return !Objects.equals(candidateRack, primaryRack);
    }
}
```

Rack Safety: Solved!

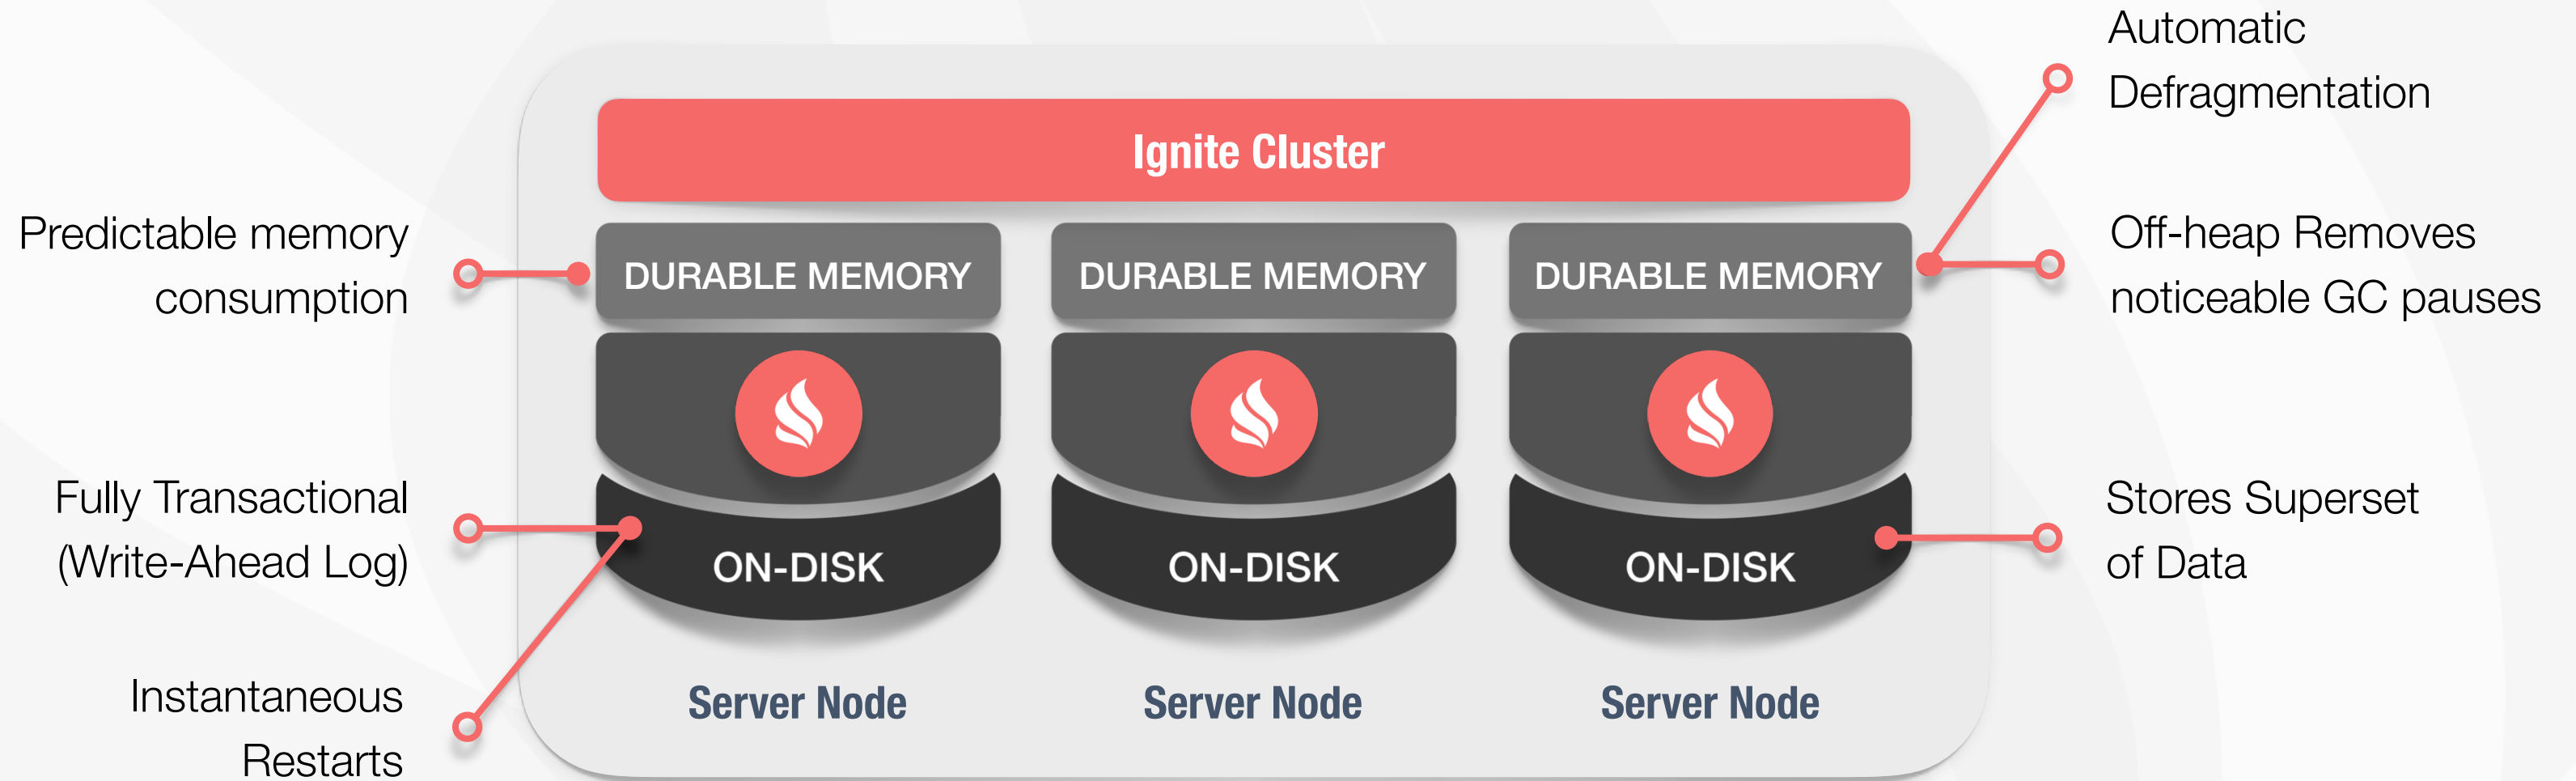


No partitions lost

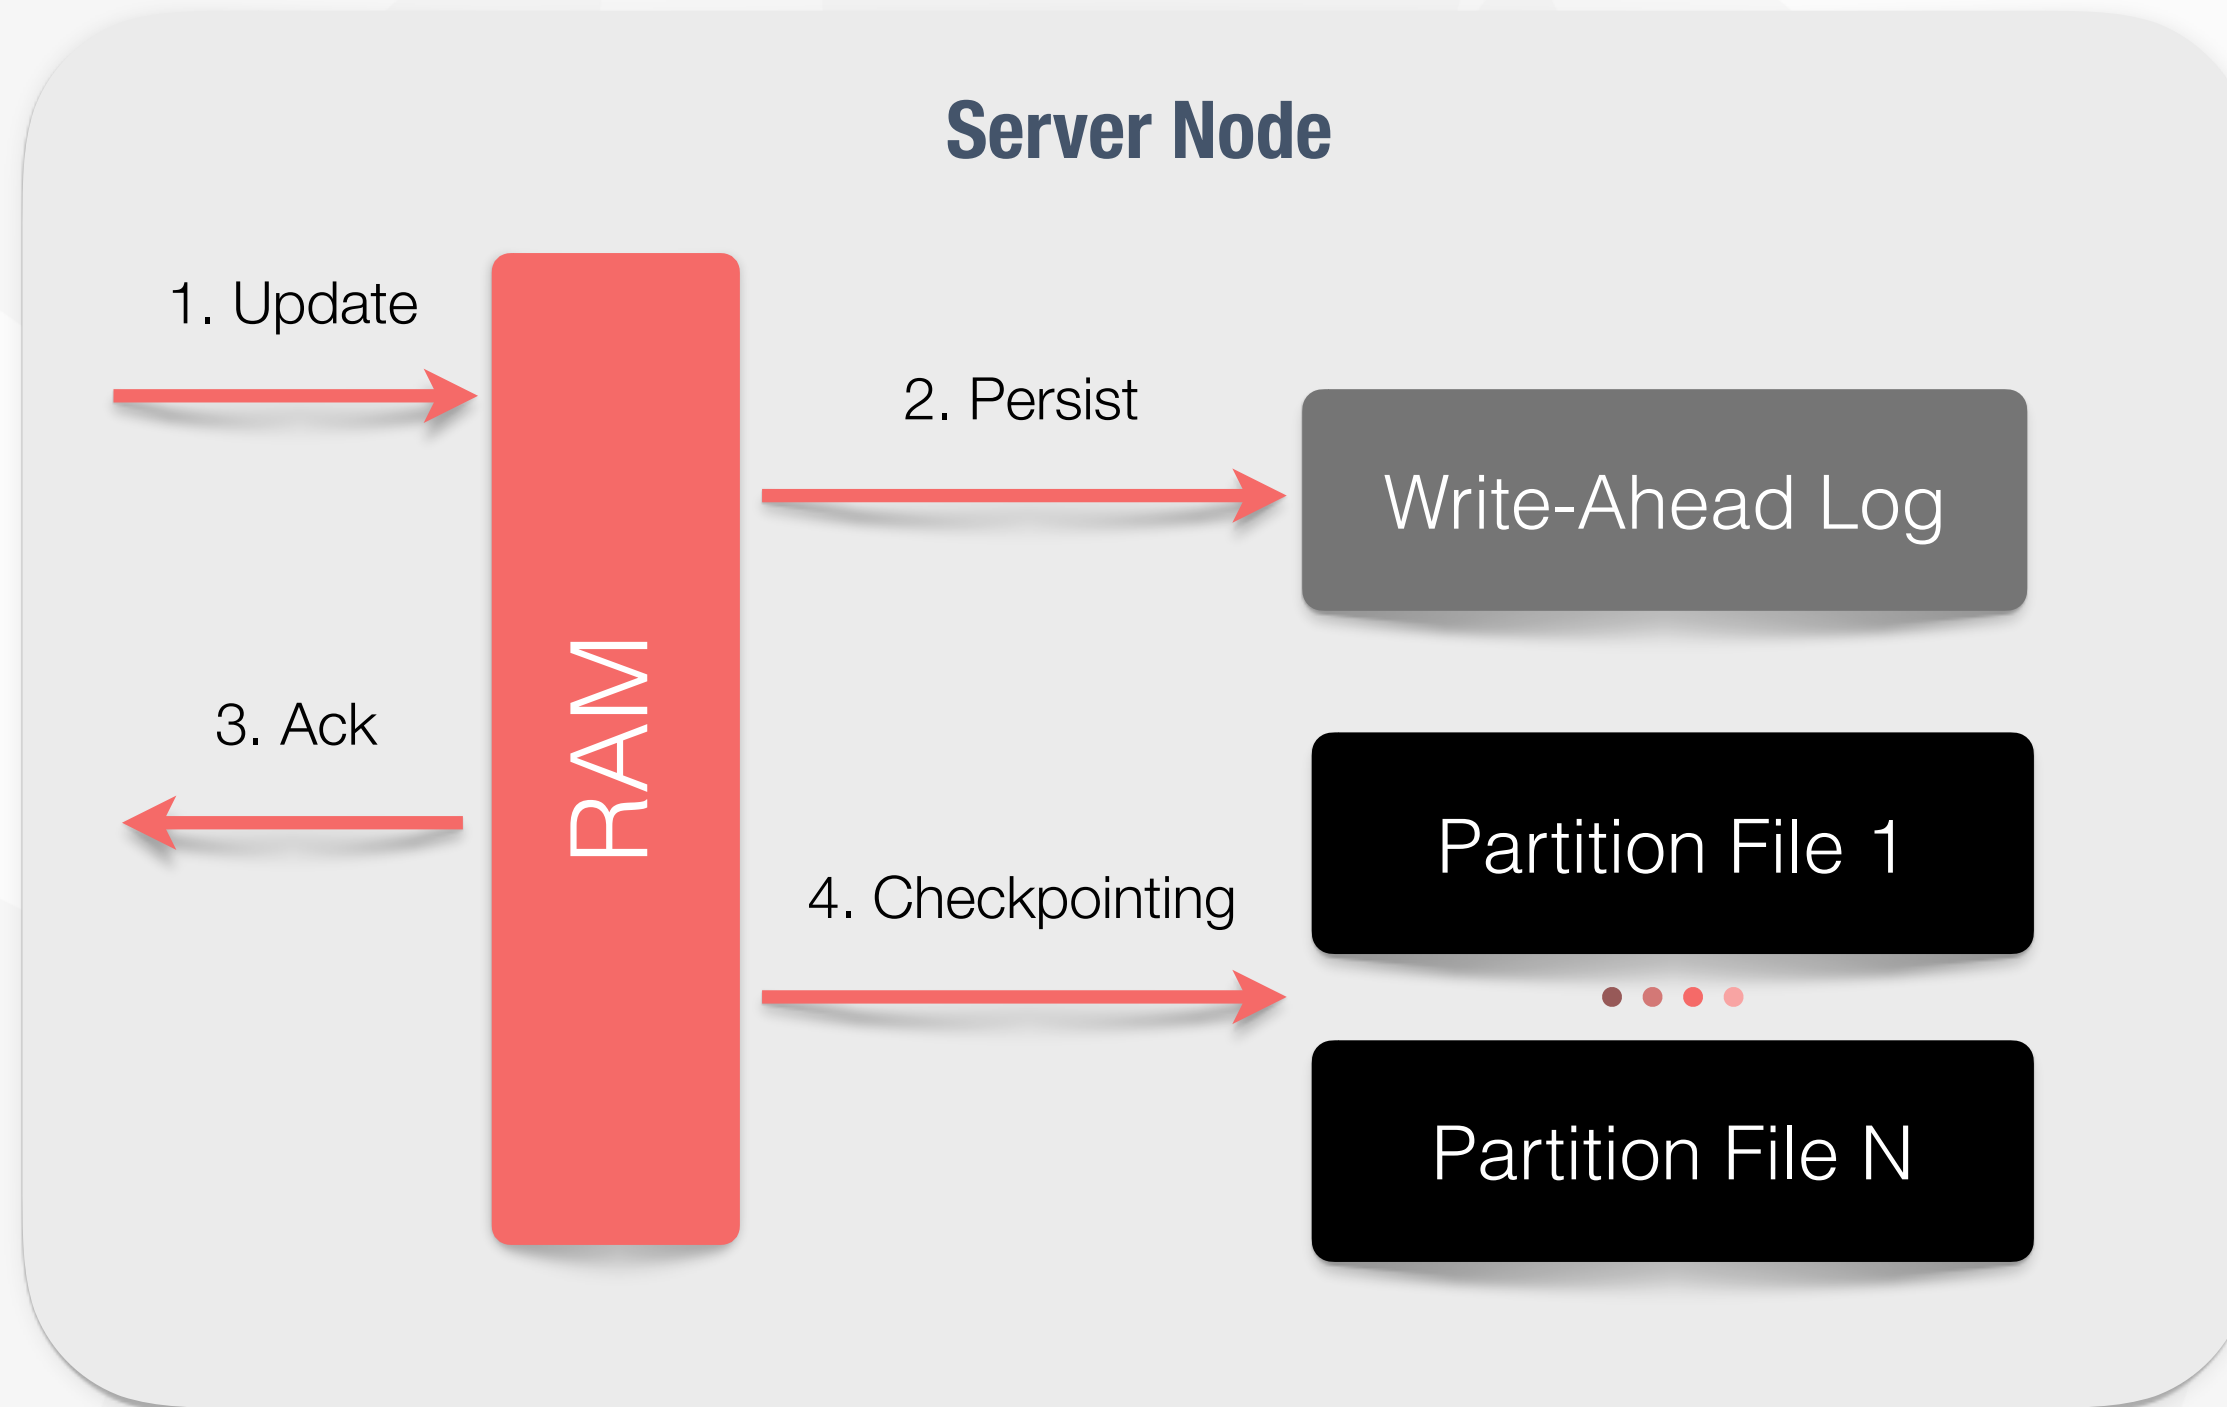
Persistence



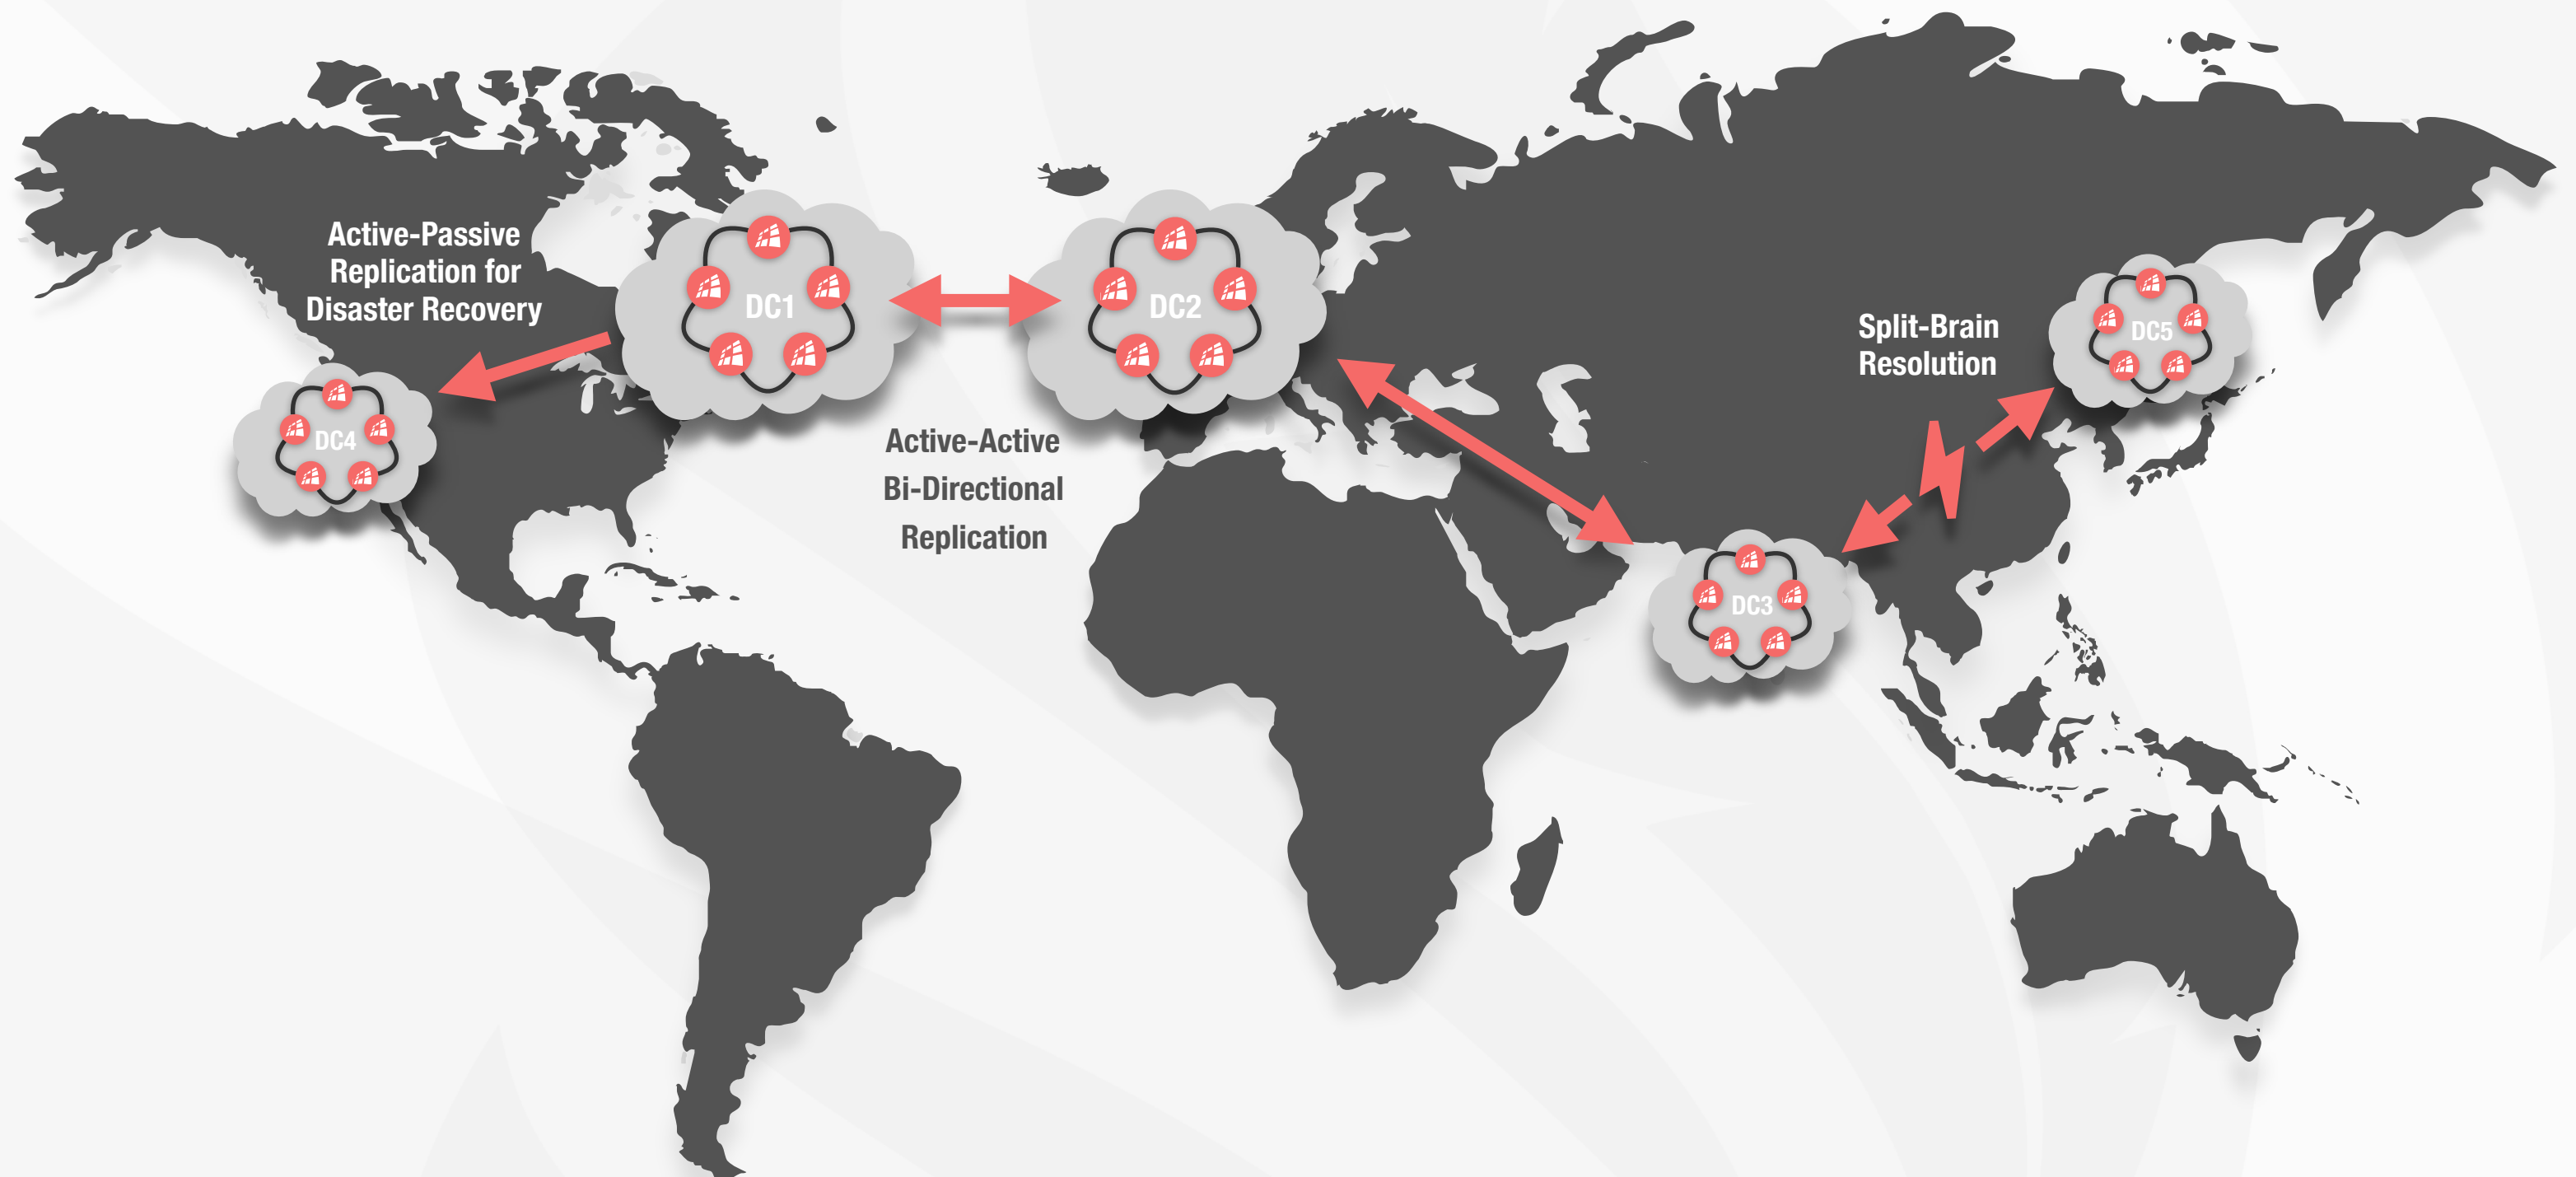
Durable Memory



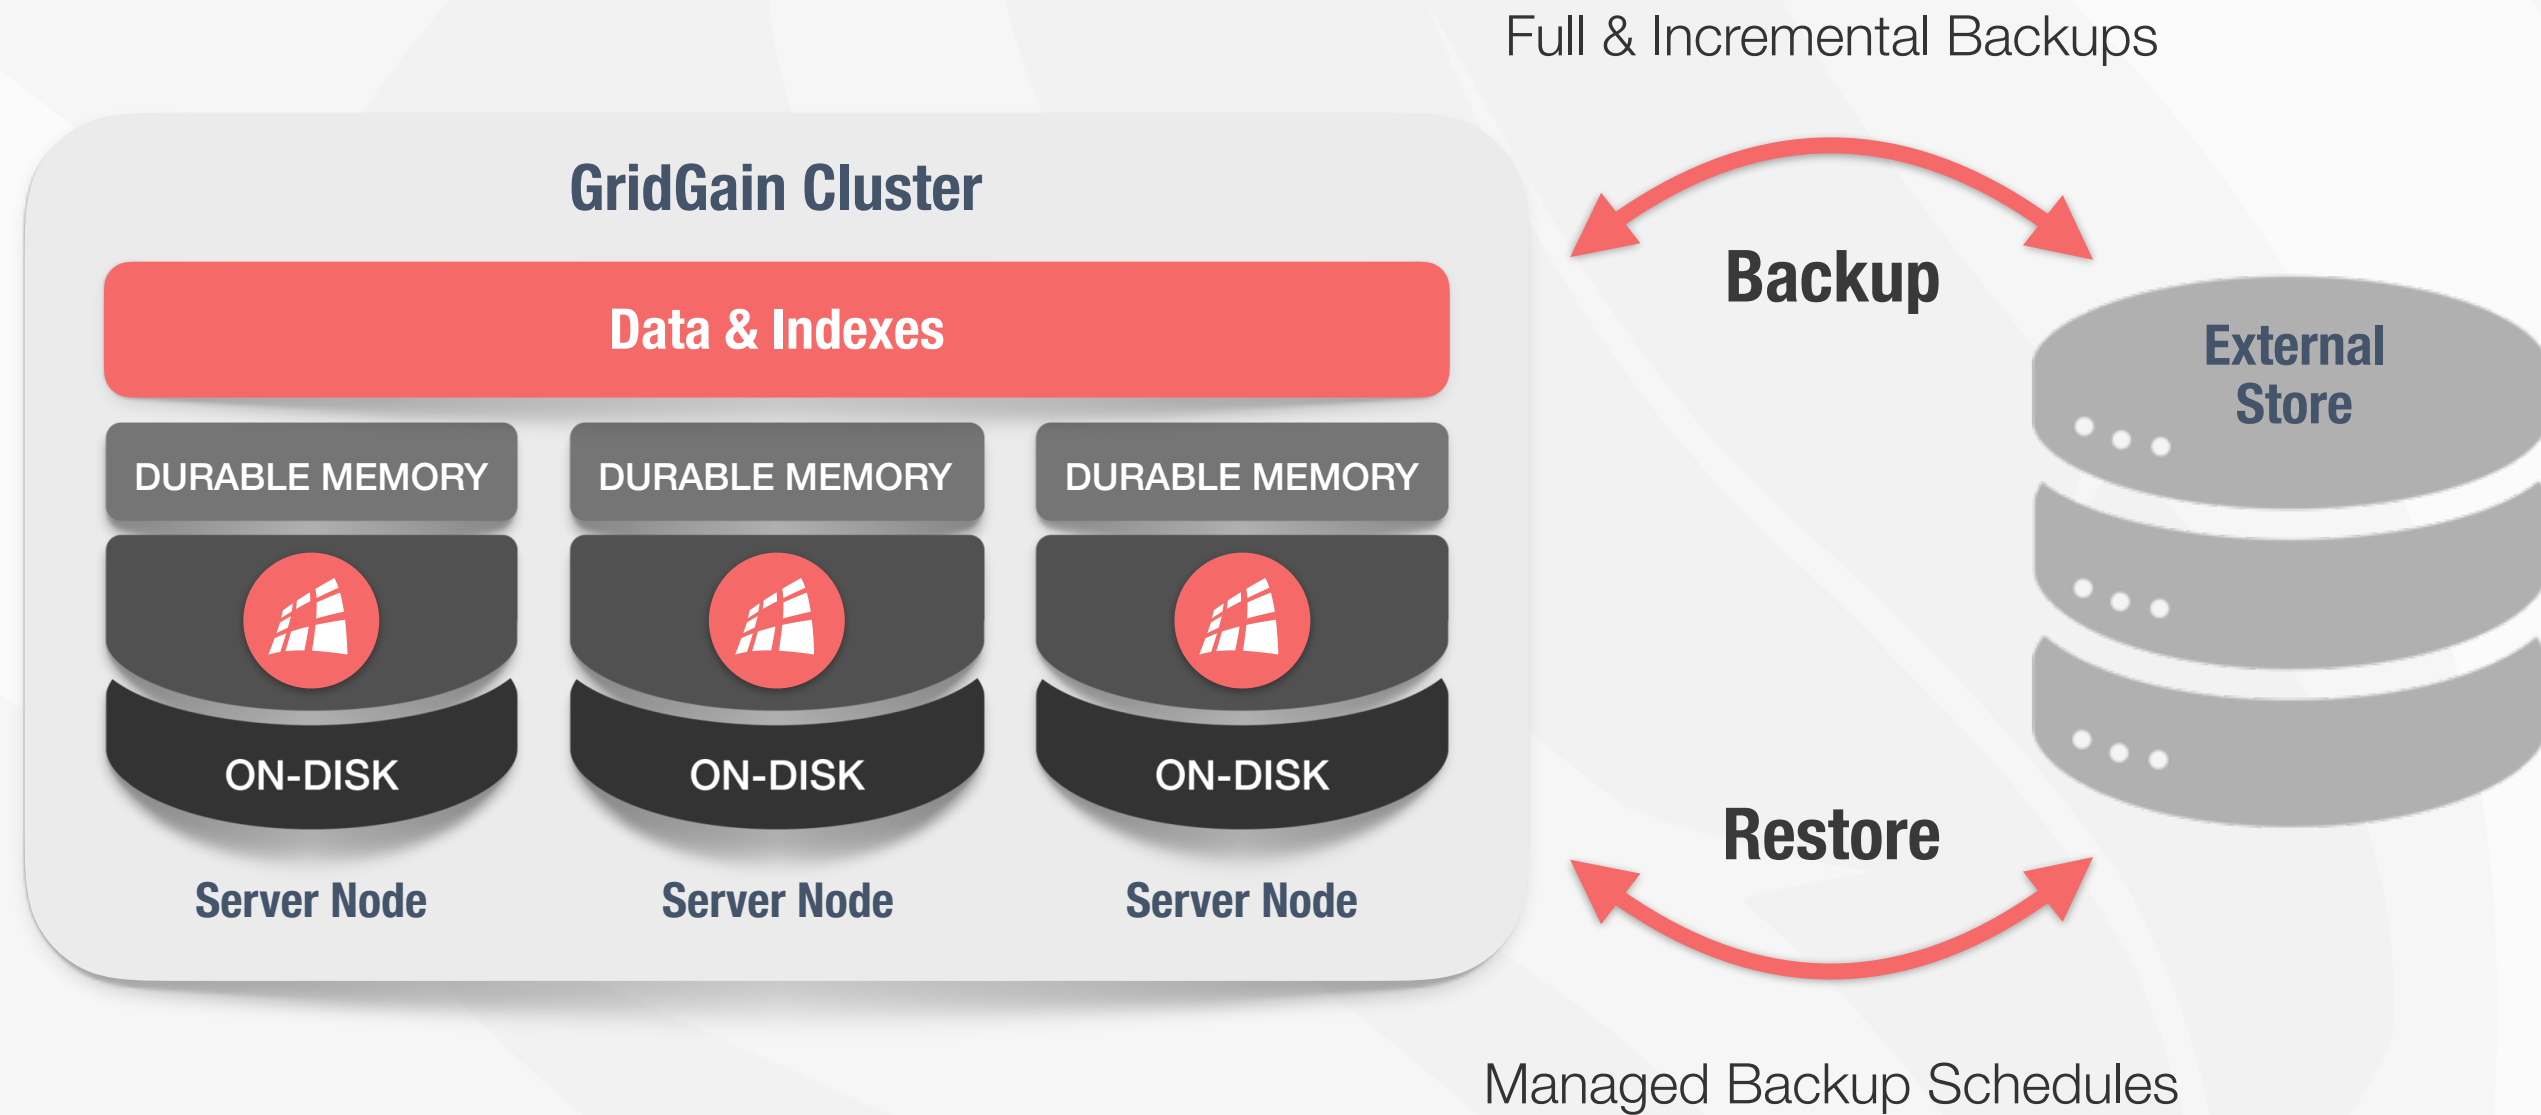
Ignite Native Persistence



GridGain Data Center Replication



Data Snapshots & Recovery





Any Questions?

Thank you for joining us. Follow the conversation.
<http://ignite.apache.org>



@apacheignite
@vkulichenko